

Chapter 1

VMware vSphere 4 Overview

VMware vSphere 4 is the new major version of VMware's flagship virtualization platform, VMware Infrastructure 3. A bare-metal hypervisor that enables full virtualization of industry-standard x86 hardware forms the foundation of this virtualization platform. In addition to this hypervisor, vSphere includes several advanced features that support innovative applications of virtualization technology. These features range from resource pooling, dynamic balancing of workloads, high availability, and disaster recovery.

VMware classifies all these vSphere features into this set of services:

- ◆ Infrastructure services
- ◆ Application services
- ◆ Management services

The infrastructure and application services are part of vSphere, and the management services are provided by VMware vCenter Server.

In this chapter, we will describe these services and provide details of their features. We will specifically call out features that are new in vSphere 4.

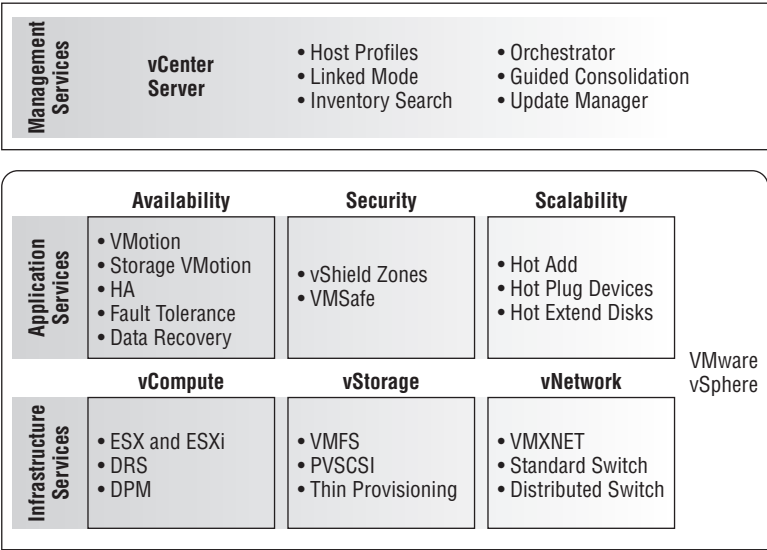
Infrastructure Services

vSphere infrastructure services are the core set of services that allows you to virtualize x86 servers (see Figure 1.1). First, these services abstract the physical x86 hardware resources, such as CPU, memory, storage, and network adapters, into virtual hardware to create virtual machines (VMs). Next, these services enable vSphere to transform resources from individual x86 servers into a shared computing platform with several operating systems and applications running simultaneously in different virtual machines. Finally, the infrastructure services provide several sophisticated features to optimize resources in such a shared environment. Figure 1.1 provides an overview of all the services in vSphere 4.

VMware vSphere provides the following types of infrastructure services:

- ◆ VMware vCompute
- ◆ VMware vStorage
- ◆ VMware vNetwork

FIGURE 1.1
VMware vSphere
overview



In the following sections, we'll provide a closer look at each type of infrastructure service.

VMware vCompute

VMware vCompute services virtualize CPU and memory resources in an x86 server. The vCompute services also aggregate these resources from several discrete servers into shared logical pools that can be allocated to applications running inside virtual machines.

The vCompute services comprise the following:

VMware ESX (and VMware ESXi) A bare-metal hypervisor that runs directly on server hardware. It supports different x86 virtualization technologies such as VMware-invented binary translation, hardware-assisted virtualization, and paravirtualization. VMware ESXi is a free version of ESX with a smaller footprint that minimizes the surface area for potential security attacks, making it more secure and reliable. ESX also includes several advanced CPU scheduling capabilities, as well as unique memory management features such as transparent page sharing and memory ballooning. These sophisticated features enable ESX to achieve higher consolidation ratios compared to its competition.

VMware Distributed Resource Scheduler (DRS) Extends the resource management features in ESX across multiple physical servers. It aggregates CPU and memory resources across many physical servers into a shared cluster and then dynamically allocates these cluster resources to virtual machines based on a set of configurable options. DRS makes sure that resource utilization is continuously balanced across different servers in the shared cluster.

VMware Distributed Power Management (DPM) Included with VMware DRS, DPM automates energy efficiency in VMware DRS clusters. It continuously optimizes server power consumption within each cluster by powering on or off vSphere servers as needed.

In the next sections, we will discuss each of these vCompute services in detail.

VMWARE ESX AND ESXi

VMware ESX and ESXi are the most widely deployed virtualization hypervisors, and they form the robust foundation of VMware vSphere. VMware ESX and ESXi use bare-metal architecture; in other words, they install directly on the server hardware, without the need for a host operating system.

VIRTUALIZATION ARCHITECTURES

Virtualization products for x86 servers commonly use two types of architectures: a hosted architecture or a hypervisor architecture. The hosted, or *type 2*, virtualization products run on top of the Windows or Linux host operating system. The host operating system controls the access to the physical resources, and virtual machines run as applications alongside other software on the host machine. The VMware Workstation, Fusion, and Server products are examples of hosted virtualization architecture.

Bare-metal hypervisor, or *type 1*, virtualization products run directly on top of the hardware with direct access and control of the hardware's resources. Since they have direct access to the hardware resources rather than going through a host operating system, the hypervisor products are more efficient than hosted virtualization products, and deliver greater scalability, robustness, and performance. VMware ESX and ESXi are examples of bare-metal hypervisor architecture.

Virtualization Technologies

VMware ESX and ESXi offer a choice of three virtualization technologies (Figure 1.2):

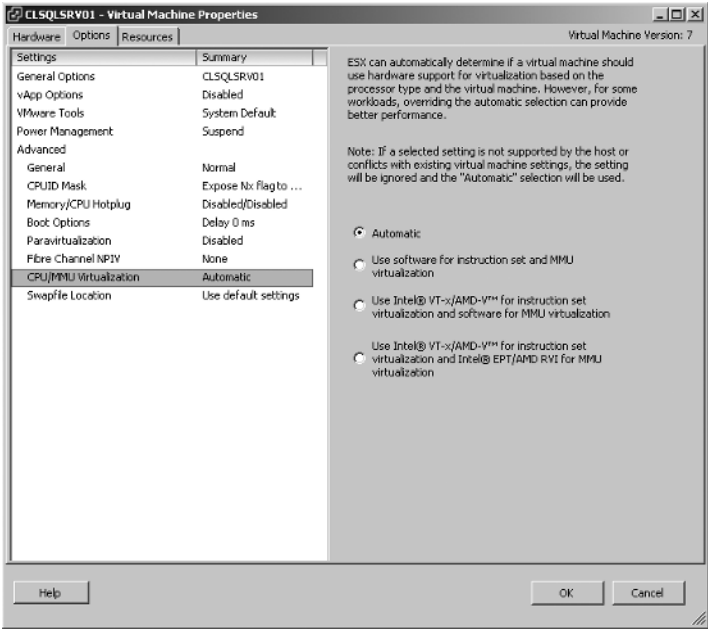
- ◆ Binary translation
- ◆ Hardware-assisted virtualization
- ◆ Paravirtualization

Binary translation is the virtualization technique that VMware invented for x86 servers. The x86 processors were not designed with virtualization in mind. These processors have 17 CPU instructions that require special privileges and can result in operating system instability when virtualized. The binary translation technique translates these privileged instructions into equivalent safe instructions, thus enabling virtualization for x86 servers. Binary translation does not require any specific features in the x86 processors and hence enables you to virtualize any x86 server in the data center without modifying guest operating system and applications running on it.

Hardware-assisted virtualization relies on the CPU instruction set and memory management virtualization features that both AMD and Intel have recently introduced in the x86 processors. The first generation of these hardware-assisted virtualization processors, called *AMD-SVM* and *Intel-VT*, only supported CPU instruction set virtualization in the processors. This alone did not perform fast enough for all different workloads, compared to the binary translation technology.

Recently, AMD and Intel have introduced newer processors that also support memory management virtualization. You'll learn more about this in the next section. Virtualization using these second-generation hardware-assisted processors usually performs better than binary translation. Consequently, with the release of vSphere, VMware ESX, and ESXi now default to hardware-assisted virtualization out of the box, but you do have the choice to override this setting.

FIGURE 1.2
VMware vSphere
virtualization
technology options



VMware vSphere also supports *paravirtualized Linux guest operating systems*—Linux kernels that include Virtual Machine Interface (VMI) support—that are virtualization-aware. Because the VMI standard is supported out of the box in newer Linux kernels, there is no need to maintain separate distributions of Linux specifically for virtualization.

Hardware-Assisted Memory Virtualization

Memory management in virtual machines is challenging compared to physical machines, especially when it comes to virtual memory address translation. In a physical machine, the operating system uses page tables to translate memory addresses from an application's "virtual" space into the machine's physical memory addresses. Similarly, in a virtual machine, guest virtual memory addresses are translated to guest physical addresses using the guest OS's page tables. However, the guest OS does not have access to the physical machine memory; ESX controls the access to the actual physical memory. ESX performs the final translation to machine physical memory addresses by implementing a set of shadow page tables for each virtual machine. Creating/maintaining the shadow page tables adds both CPU and memory overhead. This overhead can be significant for virtual machines running several processes or using multiple virtual CPUs.

Both AMD and Intel have introduced hardware-assisted memory management capabilities to alleviate this situation. Processors supporting hardware-assisted memory management implement an additional level of page tables in hardware. These hardware page tables keep track of guest

physical to machine memory address translations, which used to be maintained inside shadow page tables within ESX. Offloading this memory management to hardware has two benefits: hardware page table processing is faster than software implementation, and ESX can use the freed CPU cycles for more workload-related processing.

AMD calls its hardware-assisted memory management feature *rapid virtualization indexing* (RVI), while Intel terms its implementation *extended page tables* (EPT). ESX has supported AMD RVI since version 3.5. The support for Intel EPT was introduced in ESX 4.0.

The performance benefits of hardware-assisted memory management are achievable only if page table entries are located in hardware page tables. Remember that the real estate on a processor chip is at a premium and hence limits the size of hardware page tables. If a page table entry is not found in the hardware page tables, the associated translation lookaside buffer (TLB) miss can result in more expensive processing compared to software shadow page tables implemented by ESX. You can reduce the number of TLB misses by using large memory pages. ESX has been supporting large memory pages since version 3.5. Together, hardware-assisted memory management and large memory pages will provide better performance.

Processor Scheduling

VMware vSphere includes a sophisticated CPU scheduler that enables it to efficiently run several machines on a single ESX host. The CPU scheduler allows you to over-commit available physical CPU resources; in other words, the total number of virtual CPUs allocated across all virtual machines on a vSphere host can be more than the number of physical CPU cores available. The virtual machines are scheduled on all available physical CPUs in a vSphere host by default or can be affinitized or pinned to specific physical CPUs. The ESX CPU scheduler will also guarantee that a virtual machine only uses CPU cycles up to its configured values. When scheduling virtual CPUs allocated to virtual machines, the CPU scheduler uses a proportional-share scheduling algorithm that also takes into account user-provided resource specifications such as shares, reservations, and limits. Maintaining CPU resource allocation fairness among a number of virtual machines running on a vSphere host is a key aspect of ESX processor scheduling.

Starting with the Virtual Infrastructure 3 (VI3) release, ESX has gradually shifted from “strict” to “relaxed” co-scheduling of virtual CPUs. Strict co-scheduling required that a virtual machine would run only if all its virtual CPUs could be scheduled to run together. With relaxed co-scheduling, ESX can schedule a subset of virtual machine CPUs as needed without causing any guest operating system instability.

The ESX CPU scheduler is also aware of different processor topology architectures such as nonuniform memory access architecture (NUMA) nodes and hyperthreading.

The ESX 4.0 scheduler further improves on these capabilities by adding the following enhancements:

- ◆ More optimizations to relaxed co-scheduling of virtual CPUs, especially for SMP VMs (virtual machines with multiple virtual CPUs)
- ◆ New finer-grained locking to reduce scheduling overheads in cases where frequent scheduling decisions are needed
- ◆ Processor cache topology awareness and optimizations to account for newer processor cache architectures
- ◆ Improvements in interrupt delivery efficiency and the associated processing costs

Advanced Memory Management

VMware vSphere uses several advanced memory management features to efficiently use the physical memory available. These features make sure that in a highly consolidated environment virtual machines are allocated the required memory as needed without impacting the performance of other virtual machines. These advanced features include the following:

Memory over-commitment Similar to CPU over-commitment, memory over-commitment improves memory utilization by enabling you to configure virtual machine memory that exceeds the physical server memory. For example, the total amount of memory allocated for all virtual machines running on a vSphere host can be more than the total physical memory available on the host.

Transparent page sharing Transparent page sharing uses available physical memory more efficiently by sharing identical memory pages across multiple virtual machines on a vSphere host. For example, multiple virtual machines running Windows Server 2008 will have many identical memory pages. ESX will store a single copy of these identical memory pages in memory and create additional copies only if a memory page changes.

Memory ballooning Memory ballooning dynamically transfers memory from idle virtual machines to active ones. It puts artificial memory pressure on idle virtual machines, forcing them to use their own paging areas and release memory. This allows active virtual machines in need of memory to use this memory. Keep in mind that ESX will ensure that a virtual machine memory usage cannot exceed its configured memory.

Large memory pages Newer x86 processors support the use of large 2 MB memory pages in addition to the small 4 KB pages. Operating systems rely on the translation lookaside buffers inside the processor to translate virtual to physical memory addresses. Larger page sizes mean that a TLB cache of the same size can keep track of larger amounts of memory, thus avoiding the costly TLB misses. Enterprise applications such as database servers and Java virtual machines commonly use large memory pages to increase TLB access efficiency and improve performance. ESX supports the use of large memory pages in virtual machines and backs up with its own large memory pages to maintain efficient memory access.

Resource Management

VMware vSphere allows you to establish minimum, maximum, and proportional resource shares for CPU, memory, disk, and network bandwidth for virtual machines. The minimum resource setting or reservation guarantees the amount of CPU and memory resources for a virtual machine, while the maximum resource setting or limit caps the amount of CPU and memory resources a virtual machine can use. The proportional resource allocation mechanism provides three levels—normal, low, and high—out of the box. These settings help configure virtual machine priority for CPU and memory resources relative to each other. These can be set at the resource pool level and are inherited or overridden at the individual virtual machine level. You can leverage these resource allocation policies to improve service levels for your software applications. The key advantage of these settings is that you can change resource allocations while virtual machines are running, and the changes will take place immediately without any need to reboot.

You need to be careful when assigning the minimum settings or reservations because they guarantee resources to a virtual machine. If too much CPU and memory resources are reserved, you may not be able to start virtual machines.

New Virtual Hardware Generation

In vSphere 4.0, VMware has upgraded the virtual hardware from version 4 to version 7. This generation of virtual hardware adds the following features:

Serial attached SCSI (SAS) virtual device for Microsoft Cluster Service This virtual device is needed to support running Windows Server 2008 in a Microsoft Cluster Service configuration. Later chapters in this book cover setting up a Microsoft Cluster Service configuration using Windows Server 2008 and will demonstrate the use of this device.

IDE virtual device This virtual device is recommended for older operating systems that do not support SCSI drivers.


VMXNET generation 3 VMware introduces the third generation of their paravirtualized virtual networking adapter in vSphere 4. Refer to the “VMware vNetwork” section later in this chapter for details.

Virtual machine hot-plug support The new virtual hardware generation enables you to hot plug virtual devices to a virtual machine without having to power it off. You can hot add and remove virtual CPUs, hot add and remove network cards and disks, and hot add memory to a virtual machine when using virtual hardware version 7. The support for the CPU and memory hot-add plug-in depends upon the guest operating system support.


New virtual machines created in the vSphere environment use virtual hardware version 7 by default. However, vSphere can run virtual machines created on hosts running ESX Server versions 2.x and 3.x.

You can find the virtual hardware version either on the virtual machine Summary tab or at the top left of the Virtual Machine Properties window. To convert your existing virtual machines, you first need to update VMware Tools to the latest version. You can then upgrade the virtual hardware used by the virtual machine. Right-click the virtual machine, and the context menu should provide you with an option to perform this upgrade (Figure 1.3).

FIGURE 1.3
VM hardware
upgrade



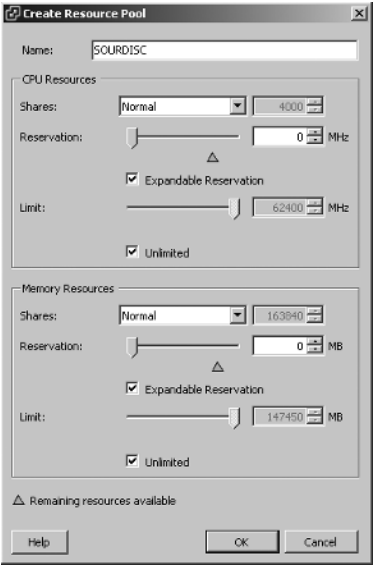
NOTE Virtual machines using virtual hardware version 7 features are not compatible with ESX/ESXi releases prior to version 4.0. The virtual machine upgrade process is irreversible, and you will see a warning window during the upgrade steps, as shown here.



DISTRIBUTED RESOURCE SCHEDULER

VMware DRS allows you to manage physical resources distributed across multiple ESX servers. Using DRS, you can aggregate CPU and memory resources from up to 32 ESX servers to create a shared pool of resources, appropriately called *resource pools*. You can then organize these resource pools to create a flexible hierarchy to reflect business priorities. DRS also allows you to extend the resource management capabilities of a single ESX server such as shares, reservations, or limits to all virtual machines within these resource pools (Figure 1.4). For example, you can assign higher shares of the total resources to the production resource pool compared to a test and development resource pool. Likewise, you can guarantee fixed CPU and memory (reservations) for business-critical applications within that production resource pool.

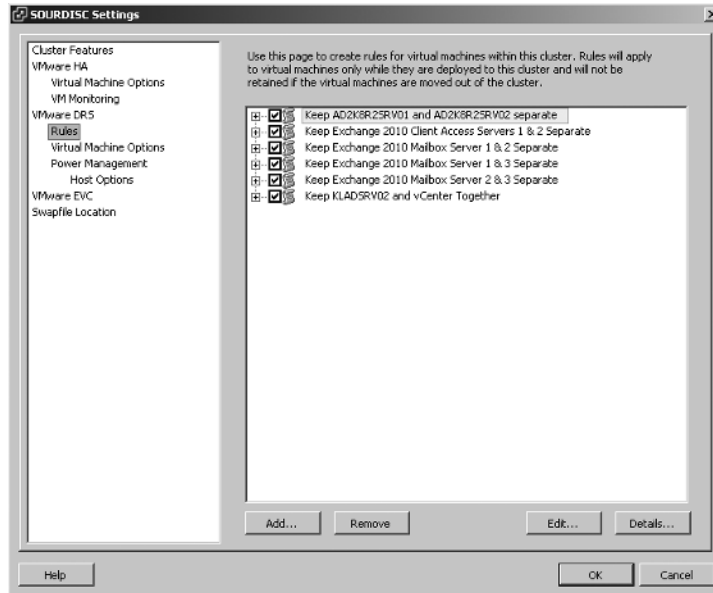
FIGURE 1.4
VMware vSphere
resource pools



DRS also allows you to define rules and policies for virtual machines' resource allocations (Figure 1.5). For example, you can define an affinity rule to make sure that all virtual machines in a given application stack always run on the same server. All network communication for such co-located virtual machines takes place in memory and can benefit application performance. Alternatively, you can define an anti-affinity rule to ensure that specific virtual machines always

run on different servers. You can use this to avoid a single point of failure and increase availability for application components such as web servers in a load-balanced farm.

FIGURE 1.5
VMware vSphere
DRS rules



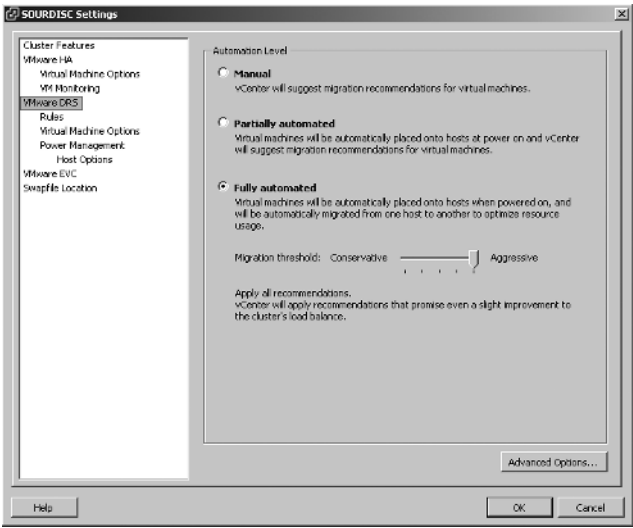
VMware DRS will help you to load balance the resource utilization across various ESX servers within resource pools. It continuously monitors utilization across resource pools and then generates recommendations to optimally place virtual machines within these resource pools. While doing this, DRS also takes into account any priority settings, such as shares, reservations, and limits, as well as resource allocation rules and policies you may have defined. You can configure to execute these virtual machine placement recommendations automatically or manually (Figure 1.6). VMware DRS does this smart resource allocation during the initial virtual machine placement and on continuous basis in real time while the virtual machine is running:

Initial placement When a virtual machine is first powered on, VMware DRS either places the virtual machine on the most appropriate ESX server automatically or makes a recommendation for you to act on.

Continuous optimization The virtual machine resource requirements change over time, and VMware DRS dynamically responds to these changing needs by migrating virtual machines between ESX servers using VMware VMotion without disrupting end users. Alternatively, you can configure DRS to use manual mode so it only makes recommendations that you can then choose to act on.

By leveraging VMotion, VMware DRS also simplifies planned maintenance on physical servers without disrupting virtual machines and end users. When you place a physical server in maintenance mode, VMware DRS identifies alternative servers where the virtual machines can run. Based on the automation mode settings, either the virtual machines are automatically moved to use the alternative servers or the system administrator performs the move manually using the VMware DRS recommendations as a guideline.

FIGURE 1.6
VMware vSphere
DRS automation
levels

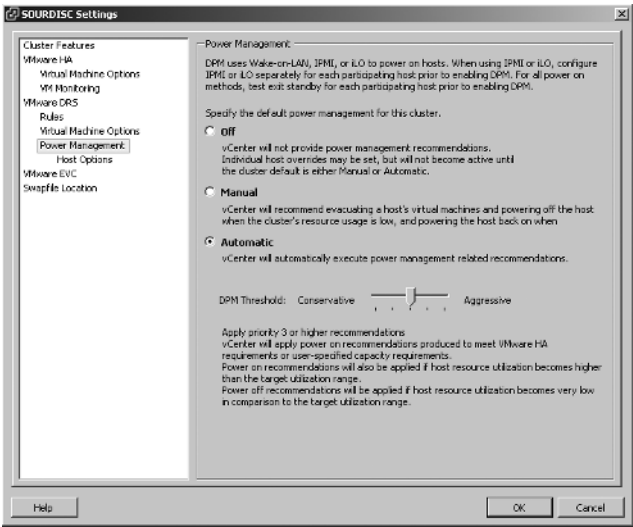


DISTRIBUTED POWER MANAGEMENT

VMware DPM enables you to reduce energy consumption in the data center by optimizing the placement of virtual machines within a cluster. When resource utilization in a DRS cluster goes down, DPM consolidates virtual machines within the cluster on fewer ESX servers and powers off unused servers to conserve energy. When resource requirements increase, DPM brings ESX servers back online to ensure that virtual machines get appropriate resources and service levels can be maintained (Figure 1.7).

DPM uses technologies such as IPMI, iLO, and Wake on LAN to control server power states. ESX/ESXi 4.0 also supports the Enhanced Intel SpeedStep and Enhanced AMD PowerNow! CPU power management technologies. These technologies leverage dynamic voltage and frequency scaling (DVFS) to reduce power consumption.

FIGURE 1.7
VMware vSphere
DPM power
management



VMware vStorage

VMware vStorage includes a number of advanced features to efficiently manage and use storage in virtualized environments while trying to hide the complexities of the underlying storage systems. The key features include the following:

- ◆ VMware vStorage *Virtual Machine File System* (VMFS) is a custom-designed cluster file system specifically optimized to store and manage virtual machines. It allows the efficient sharing of back-end storage by multiple ESX servers and is the key enabler for VMware features such as VMotion, Storage VMotion, DRS, VMware High Availability, and Fault Tolerance.
- ◆ VMware vStorage *thin provisioning* enables you to over-commit storage capacity similar to CPU and RAM over-commitment. Thin provisioning initially allocates only as much storage capacity as needed and grows as more data is stored in the virtual disk. This increases storage utilization allowing you to defer storage purchases until they are really required. This can significantly reduce an organization's storage budget.
- ◆ vSphere introduces a high-performance, *paravirtualized SCSI (PVSCSI) storage adapter* that offers greater throughput and lower CPU utilization for virtual machines. This is best suited for environments that run very I/O-intensive guest applications.
- ◆ VMware vStorage offers APIs for data protection, multipathing, and storage replication adapters so that storage partners can integrate their solutions with vSphere.

In the next sections, we will discuss each of these vStorage services in detail.

VIRTUAL MACHINE FILE SYSTEM

VMware created a cluster file system, VMFS, specially designed to store and manage virtual machines. VMFS is optimized to support large files associated with virtual disks, thus enabling encapsulation of an entire virtual machine in a set of files. Using VMFS, you can place these virtual machine files on a shared storage and allow multiple ESX servers to concurrently read and write to this shared storage.

By managing concurrent access to the shared back-end storage, VMFS enables the foundation for key VMware features such as VMotion, Storage VMotion, DRS, VMware High Availability, and Fault Tolerance. As virtual machines are migrated to or restarted on different ESX servers, VMFS ensures that individual ESX servers are not single points of failure and helps DRS to balance resource utilization across multiple servers. VMFS uses on-disk file locking to ensure that the same virtual machine is not powered on by multiple servers at the same time.

VMFS also acts a logical volume manager by providing an interface to different types of storage such as Fibre Channel SAN, iSCSI SAN, and NAS. VMFS hides the complexities of underlying storage systems and, irrespective of the storage type, simplifies storage management using automatic discovery and mapping of LUNs to a VMFS volume. You can connect or disconnect a VMware ESX server from a VMFS volume without impacting other VMware ESX hosts. vSphere also adds dynamic growth capabilities to VMFS without the need for any downtime. These new capabilities include hot expansion of VMFS volumes and virtual disks stored in VMFS.

vStorage VMFS Volume Grow The VMFS Volume Grow capability in vSphere allows you to dynamically expand the size of an existing data store that resides on a VMFS volume without disrupting running virtual machines. It complements the dynamic LUN expansion capability that exists in many storage array offerings today. After you expand the LUN where a data store resides through an array management utility, you can use VMFS Volume Grow to expand the VMFS extent on the expanded LUN. You can verify the increased VMFS volume (data store) size from vCenter Server.

For earlier versions of ESX, you have to use VMFS spanning across multiple LUNs to increase the size of an existing VMFS volume. First, you expand the LUN upon which the VMFS volume resides; next, you create a separate disk partition in that additional storage space and add the new partition as if you were adding a second LUN to the VMFS volume.

Hot extend for virtual disks Hot extend for virtual disks allows you to add virtual storage to running virtual machines without any downtime. You can use hot extend for VMFS flat virtual disks using persistent mode and for ones that do not have any VMFS snapshots. You will need to run guest operating system tools for it to start using the additional storage. Together with the VMFS Volume Grow capability, this feature provides a very flexible and dynamic way to manage storage capacity growth.

VIRTUAL DISK THIN PROVISIONING

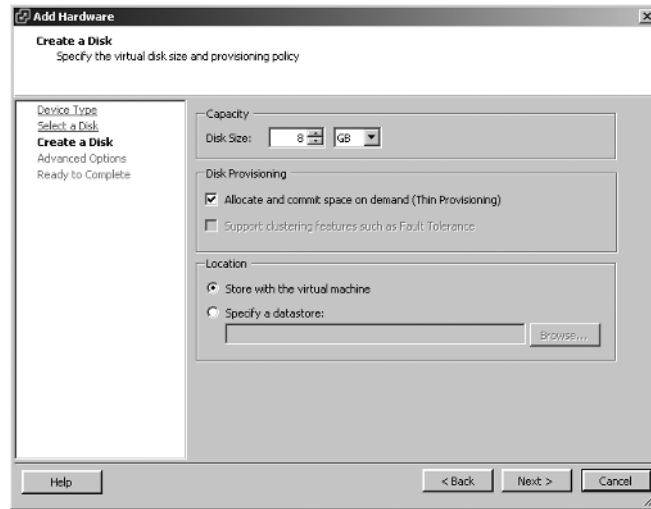
VMware virtual disk thin provisioning enables over-commitment of storage capacity similar to CPU and RAM over-commitment. It allows the storage administrator to configure more virtual machine storage capacity than the actual physical storage currently available. This is possible because thin provisioning enables virtual machines to utilize storage space on an as-needed basis. When a virtual disk is initially allocated, it is assigned 1 MB of storage space in the data store. As that disk space is used up, additional 1 MB chunks of storage are allocated for the virtual disk so that the underlying storage demand will grow as its size increases. This dynamic allocation reduces storage over-provisioning and increases storage utilization, allowing you to defer storage purchases until they are really required. This can significantly reduce an organization's storage budget.

VMware's desktop products such as VMware Workstation have long provided the feature to allow virtual machines to allocate storage space as needed. In Virtual Infrastructure 3, thin provisioning was used by default for virtual disks created on NFS data stores and was available for block-based data stores through the command line. With VMware vSphere, vCenter now fully supports virtual disk thin provisioning for all virtual disks when deploying or migrating virtual machines (Figure 1.8).

vCenter Server 4.0 continuously monitors the storage allocations and tracks it against storage capacity so that it can generate alerts and alarms to warn vSphere administrators against any future "out of space" situations.

NOTE Virtual disk thin provisioning should not be confused with the same thin provisioning that an array vendor might offer. In fact, with vSphere, you now have the capability of doing thin provisioning at the data store level in addition to doing thin provisioning at the storage array.

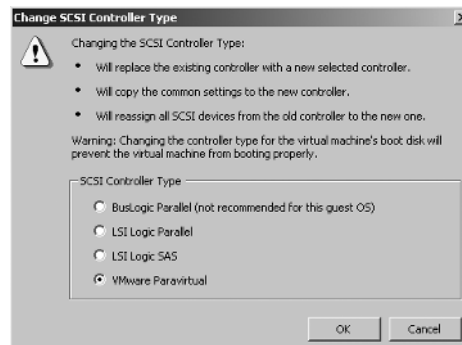
FIGURE 1.8
Creating a thin-
provisioned virtual
hard disk



VMWARE PARAVIRTUALIZED SCSI

Early versions of ESX supported virtual SCSI adapters that emulated BusLogic and LSI Logic hardware storage adapters. This full virtualization provided broad compatibility with guest operating systems supported by ESX. However, this prevents ESX from taking advantage of performance optimizations that can be achieved with paravirtualized devices such as VMXNET network adapters. vSphere introduces a new virtual storage adapter called PVSCSI, as shown in Figure 1.9, which extends to the storage stack performance gains typically associated with other paravirtual devices. In that respect, the PVSCSI adapter is similar to the paravirtualized network adapter VMXNET that is available in ESX. As with other paravirtual devices, the PVSCSI adapter improves I/O efficiency by using optimizations such as batching the processing of I/O requests and I/O completion interrupts and reducing the cost of virtual interrupts. The PVSCSI adapter also benefits from a reduced number of context switches between the guest operating system and ESX virtual machine monitor.

FIGURE 1.9
Paravirtualized
SCSI controller



The performance benefits of the PVSCSI driver are visible for virtual machines issuing more than 2,000 I/O requests per second. For lower I/O throughput, VMware recommends that you continue to use the default or primary adapter. For example, LSI Logic is the default primary adapter for virtual machines with Microsoft Windows 2008 guest operating systems. You can use it for the virtual disk that hosts the system software (boot disk) and a separate PVSCSI adapter for the disk that stores user data, such as a database.

In vSphere 4.0, you cannot use the PVSCI adapter for a boot partition, but subsequent versions are expected to support this. The PVSCSI driver currently works with only certain guest OS versions such as Windows Server 2003, Windows Server 2008, and RHEL 5. It can also be shared by multiple virtual machines running on a single ESX, unlike the VMDirectPath I/O, which will dedicate a single adapter to a single virtual machine.

vSTORAGE APIs

vSphere introduces several storage APIs to enable integration with leading storage hardware and software products for data protection, high availability, and disaster recovery. Storage partners have written plug-ins to vSphere using a pluggable storage architecture, which is an open modular framework to leverage vendor-specific capabilities for better performance. You can use these partner plug-ins for better flexibility and visibility when configuring storage resources for your deployment. The vStorage APIs include the following:

The vStorage APIs for Multipathing These provide an I/O multipathing framework for storage partners to create Multipathing Extensions Modules that plug in to VMware ESX/ESXi to deliver storage path failover and storage I/O throughput optimized for partners' storage arrays.

The vStorage API for Data Protection This enables backup tools to directly connect the ESX servers and the virtual machines running on them without any additional software installation. This API allows backup tools to do efficient incremental, differential, and full-image backups and restores of virtual machines. To avoid any service disruption, this API also makes it possible to offload backup processing from ESX servers.

STORAGE PERFORMANCE AND SCALABILITY

VMware has made significant performance and scalability improvements to the storage stack in vSphere. These enhancements apply to all supported storage protocols: Fibre Channel SAN, iSCSI SAN, and NFS. Together with the new paravirtualized SCSI driver, these storage stack optimizations dramatically improve storage I/O performance, in terms of both I/O throughput and the CPU cost per I/O. These are the key storage performance improvements:

Improved I/O efficiency With VMware vSphere 4, like earlier versions of ESX, you can achieve I/O throughput levels that are limited only by the capabilities of the underlying storage system and the storage connectivity link speed. Because of storage stack optimizations, vSphere 4 uses fewer CPU cycles to achieve these throughput levels.

Software iSCSI and NFS support with jumbo frames Using jumbo frames is a recommended best practice to improve performance for Ethernet-based storage. Earlier ESX versions supported jumbo frames only for networking. In vSphere, you can now leverage jumbo frames for both NFS and iSCSI storage whether you use 1 Gbps or 10 Gbps NICs.

iSCSI support improvements In vSphere 4, VMware has rewritten the entire iSCSI software initiator stack for both software iSCSI (that is, in which the iSCSI initiator runs at the ESX layer) and hardware iSCSI (that is, in which ESX leverages a hardware-optimized iSCSI HBA). As a result, both software and hardware iSCSI in vSphere 4 provide better throughput and CPU efficiency when compared to the earlier 3.5 version.

VMDirectPath I/O for Storage

VMDirectPath I/O is a new capability in vSphere that enables virtual machines to directly access the underlying physical I/O devices. When using VMDirectPath I/O for storage, there is a one-to-one mapping between an HBA and a VM, and you are not allowed to share an HBA by more than one VM. VMDirectPath is designed to handle the I/O needs of special-purpose I/O appliances and I/O-intensive virtual machines. By accessing the I/O devices directly and bypassing the hypervisor, the guest OS enhances CPU efficiency in handling the high I/O workloads. However, this I/O throughput scalability comes at the cost of other virtualization features. Features such as VMotion, hardware independence, and sharing of physical I/O devices are not available for virtual machines using VMDirectPath I/O.

In vSphere 4.0, VMDirectPath I/O is experimentally supported for the following storage I/O devices:

- ◆ QLogic QLA25xx 8 Gb Fibre Channel adapters
- ◆ Emulex LPe12000 8 Gb Fibre Channel adapters
- ◆ LSI 3442e-R and 3801e (1068 chip based) 3 Gb SAS adapters

VMware vNetwork

VMware vNetwork provides features to help you deploy and manage enterprise-class virtual networking that can communicate with external physical networks. It comprises the following key features:

- ◆ *Virtual network adapters* enable network virtualization by allowing virtual machines to network like physical machines do. VMware provides three types of virtual network adapters that virtual machines can use. vSphere introduces the third generation of the paravirtualized in-guest network drivers (VMXNET3) for enhanced network I/O performance.
- ◆ The VMware vNetwork *Standard Switch* enables you to create a virtual network between virtual machines within a single VMware ESX/ESXi host as well as those on the outside physical network. These virtual switches support the same networking protocols as physical switches and implement enterprise-class features such as VLANs and hardware NIC teaming for availability and performance.
- ◆ The VMware vNetwork *Distributed Switch* moves beyond per-host network configuration and simplifies networking management across multiple hosts in VMware vSphere environments from one central interface. It also enables third-party distributed virtual switches such as the Cisco Nexus 1000V Series virtual switch to be used in VMware vSphere environments so that network administrators can use familiar interfaces when managing virtual networking.

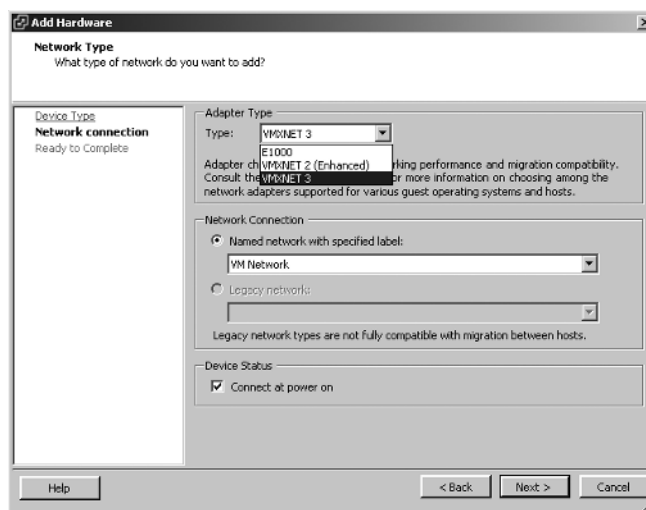
We will discuss each of these vNetwork services in detail in the next sections.

VIRTUAL NETWORK ADAPTERS

VMware supports three types of virtual network adapters for virtual machines. The AMD Lance PCNet32 and Intel E1000 physical network adapters provide full virtualization that is compatible with most guest operating systems. VMware also provides the high-performance paravirtualized network adapter VMXNET for networking virtual machines. In vSphere 4, VMware introduces VMXNET 3, the third generation of this paravirtualized network adapter (Figure 1.10). VMXNET3 builds upon the earlier versions, VMXNET and Enhanced VMXNET, by adding these features:

- ◆ MSI/MSI-X support (subject to guest operating system kernel support)
- ◆ Receive-side scaling (supported in Windows 2008 when explicitly enabled through the device's Advanced Configuration tab)
- ◆ IPv6 checksum and TCP segmentation offloading (TSO) over IPv6
- ◆ VLAN offloading
- ◆ Large TX/RX ring sizes (configured from within the virtual machine)

FIGURE 1.10
Setting a VMX-
NET3 network
adapter



VMware includes two other network adapters, vswif and vmknic, for ESX/ESXi Service Console and VMkernel, respectively. All these virtual network adapters support both IPv4 and IPv6. The support for IPv6 for the ESX/ESXi VMkernel and Service Console was added in vSphere 4. vCenter Server 4 also supports IPv6 so that vSphere customers can manage mixed IPv4/IPv6 environments. IPv6 support for network storage is considered experimental in vSphere 4.0.

vNETWORK STANDARD SWITCH

Virtual switches allow virtual machines on the same ESX Server host to communicate with each other using virtual ports and with an external network using uplink ports. These switches do

not need any additional networking hardware and are implemented in software based on the configuration you specify. Virtual switches are similar to physical switches in several ways:

- ◆ They use same networking protocols.
- ◆ They support VLANs compatible with standard VLAN implementations.
- ◆ They support Layer 2 forwarding.
- ◆ They support offloading features for TCP checksum and segmentation.
- ◆ They support Layer 2 security policies such as locking down MAC address changes.

The VMware virtual switch is also capable of binding multiple virtual NICs together, similar to NIC teaming in physical servers, to offer high availability and throughput for the virtual machines.

With more than 1,000 virtual ports per switch, you can support a large number of virtual machines per single virtual switch. Each virtual switch is isolated, and you cannot connect multiple virtual switches within the same vSphere host. This helps improve security for virtual networks, in addition to the Layer 2 security features listed earlier.

You can compare the features available in VMware virtual switches at www.vmware.com/products/vnetwork-distributed-switch/features.html.

vNETWORK DISTRIBUTED SWITCH

With vSphere, VMware has introduced the vNetwork Distributed Switch (vDS) that aggregates virtual switches from multiple vSphere hosts in a single distributed network switch. This eliminates the restriction of managing virtual networking for each vSphere host separately, and it simplifies network maintenance for the entire vSphere cluster. vDS provides a centralized interface from VMware vCenter Server for provisioning, administration, and monitoring of virtual networking for your entire data center. This can significantly reduce ongoing network maintenance activities and allow you to quickly scale up networking capacity. vDS also enables the following features:

- ◆ Network VMotion
- ◆ Bidirectional traffic shaping
- ◆ Third-party virtual switch support with the Cisco Nexus 1000V Series virtual switch

Network VMotion

Network VMotion maintains virtual machine networking state (for example, counters and port statistics) as the virtual machine moves from one host to another on a vDS. As a result, you have a consistent view for the virtual network interface regardless of which vSphere host a virtual machine is located on or how frequent a virtual machine is migrated by VMotion. This is very helpful in monitoring and troubleshooting any network-related activities in large-scale vSphere deployments.

Bidirectional Traffic Shaping

vNetwork Standard Switches allow you to set up traffic shaping for egress or transmit (from virtual machine to network) traffic. vDS expands this capability to include bidirectional traffic shaping. Both egress (from virtual machine to network) and ingress (from network into virtual machine) traffic-shaping policies can now be applied on DV port group definitions. You can use the following three characteristics to define traffic-shaping policies:

- ◆ Average bandwidth
- ◆ Peak bandwidth
- ◆ Burst size

You can leverage traffic shaping to limit the traffic to or from a virtual machine or group of virtual machines to protect either a virtual machine or other traffic in an over-subscribed network.

Third-Party Virtual Switch Support with the Cisco Nexus 1000V Series Virtual Switch

The vDS includes support for third-party distributed virtual switches. Cisco collaborated with VMware to leverage this extensibility to develop the Cisco Nexus 1000V Series virtual switch.

Both the Cisco Nexus 1000V Series virtual switch and the VMware vNetwork Distributed Switch use the same distributed switching model.

Virtual Ethernet modules (VEMs) Each ESX host implements VEMs, which are the switching data planes, and provide the frame-forwarding capabilities. These VEMs leverage the ESX host APIs and can support the same physical NICs and hardware compatibility list (HCL) as the VMware Standard Switch and vNetwork Distributed Switch.

Virtual supervisor modules (VSMs) The Cisco NX-OS operating system implements VSMs. They provide the control plane function for the VEMs and can exist as a guest VM or stand-alone appliance. VSMs allow you to use the familiar Cisco command-line interface (CLI) for management and configuration. You can also use vSphere Client to communicate with VSMs and vCenter Server for optional management and configuration.

The Cisco Nexus 1000V offers an expanded feature set compared to the VMware vNetwork Distributed Switch and is similar to that provided by the physical Cisco Catalyst and Nexus switches. You can find more information on the Cisco Nexus 1000V at <http://cisco.com/go/nexus1000v>.

NETWORK PERFORMANCE AND SCALABILITY

vSphere includes several performance enhancements to the network stack. These enhancements come in two ways: support for various offload technologies and optimizations to the existing network stack processing. It is a common performance optimization practice to offload various network processing tasks to reduce the CPU overhead associated with processing network I/O. The VMXNET3 network adapter supports performance offload technologies such as TCP Segmentation Offloading, TCP/IP Checksum Offload, and Large Receive Offload, as well as other optimizations like jumbo frames. vSphere 4 also includes optimizations to the network

stack, such as NetQueue, that can saturate even 10 Gbps links for both transmit and receive-side network I/O. You will also notice a significant increase in iSCSI throughput and maximum network throughput for VMotion because of the VMkernel TCP/IP stack optimizations in vSphere 4.

VMDIRECTPATH I/O FOR NETWORKING

VMDirectPath I/O for Networking is a new capability in vSphere that enables virtual machines to directly access underlying physical I/O devices. Similar to VMDirectPath I/O for Storage, VMDirectPath is designed for special-purpose I/O appliances and network I/O-intensive virtual machines that require very efficient network stack processing for high throughput but do not need to support additional virtualization features such as VMotion, Fault Tolerance, and suspend/resume.

In vSphere 4, VMDirectPath I/O for Networking is supported for the following devices:

- ◆ Intel 82598 10 Gigabit Ethernet Controller
- ◆ Broadcom 57710 and 57711 10 Gigabit Ethernet Controller

Application Services

Virtualization enables interesting use cases such as VMotion and VMware HA that can be used by any application running inside virtual machines. VMware vSphere application services build upon such use cases to deliver enterprise readiness features for applications running inside VMware virtual machines. These services are expected to enhance the service levels for virtualized applications more easily compared to physical deployments.

VMware vSphere provides the following types of application services:

- ◆ Availability
- ◆ Security
- ◆ Scalability

We'll take a closer look at each type of application service in the next sections.

Availability

Improving availability for applications is probably the most innovative and exciting use of virtualization technology. With availability services in vSphere, you can lower both planned and unplanned downtime for all applications running inside VMware virtual machines. Furthermore, vSphere enables this high availability without the need for complex hardware or software clustering solutions.

To minimize service disruptions because of *planned* hardware downtime, VMware vSphere includes the following availability services:

VMware VMotion Using VMware VMotion, you can migrate running virtual machines from one vSphere server to another without impacting the applications running inside virtual machines. The end users do not experience any loss of service. You can leverage VMotion to

move virtual machines off a vSphere server for any scheduled hardware maintenance without the need for any application downtime.

VMware Storage VMotion VMware Storage VMotion enables similar functionality at the storage level. You can migrate virtual disks of running virtual machines from one storage array to another with no disruption or downtime. Storage VMotion will help you avoid any application downtime because of planned storage maintenance or during storage migrations.

To reduce service disruptions because of *unplanned* hardware downtime, VMware vSphere availability services include the following features:

VMware High Availability (HA) If a virtual machine goes down because of hardware or operating system failures, VMware HA automatically restarts the virtual machine on another ESX server within minutes of the event. VMware HA provides a much simpler and cost-effective high-availability solution compared to traditional clustering solutions.

VMware Fault Tolerance (FT) VMware FT improves high availability beyond VMware HA. By maintaining a shadow instance of a virtual machine and allowing immediate failover between the two instances, VMware FT avoids even the virtual machine reboot time required in the case of VMware HA. Thus, it prevents any data loss or downtime even if server hardware fails. Like VMware HA, it can be a cheaper and simpler alternative to traditional clustering solutions.

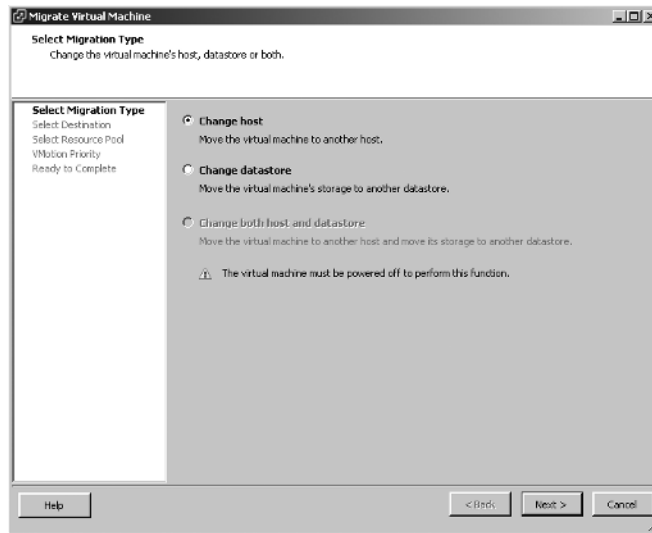
VMware Data Recovery VMware Data Recovery enables a simple disk-based backup and restore solution for all of your virtual machines. It does not require you to install any agents inside virtual machines and is completely integrated into VMware vCenter Server. VMware Data Recovery leverages data deduplication technology to avoid saving duplicate storage blocks twice, thus saving both backup time and disk space.

VMWARE VMOTION

VMotion enables live migration of running virtual machines from one ESX server to another with no downtime (Figure 1.11). This allows you to perform hardware maintenance without any disruption of business operations. The migration of the virtual machine is quite seamless and transparent to the end user. When you initiate a VMotion, the current state of the virtual machine, along with its active memory, is quickly transferred from one ESX server to another over a dedicated network link, and the ESX server gets the control of virtual machine's storage using VMFS. The virtual machine retains the same IP address after the migration. VMotion is the key enabling technology that allows VMware DRS to create a self-managing, highly optimized, and efficient virtual environment with built-in load balancing.

vSphere supports Enhanced VMotion Compatibility (EVC) using CPU features such as Intel FlexMigration and AMD-V Extended Migration technologies to allow migrations from older servers to newer servers.

FIGURE 1.11
VMware vSphere
VMotion types



VMWARE STORAGE VMOTION

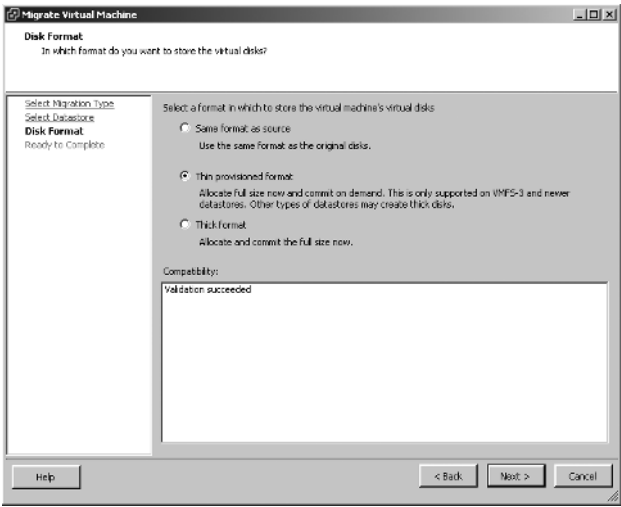
Storage VMotion allows you to migrate virtual machine disks for running virtual machines from one storage array to another. This avoids the need to schedule any service disruption or downtime because of planned storage maintenance or during storage migrations and upgrades. Storage VMotion was introduced in VMware Infrastructure 3.5, but it supported only the CLI and a limited number of storage protocols. Storage VMotion in vSphere is now integrated into vCenter and features several new capabilities:

- ◆ Includes full support for migrations from one storage vendor to another
- ◆ Works across NFS, Fibre Channel, and iSCSI storage protocols
- ◆ Includes an option to change the virtual disk format during a Storage VMotion session such as changing a thick to thin virtual disk format, or vice versa.

Storage VMotion in vSphere also features a more efficient migration process. It leverages a new and more efficient block copy mechanism in the VMkernel instead of using the virtual disk snapshot method found in the earlier releases of VMware products.

You can use Storage VMotion with thin provisioning to not only migrate a VM from one data store to another but also reclaim over-provisioned storage space during this process. Virtual machines often have guest OS virtual disks that are over-provisioned compared to what they truly need for their current storage requirements. With Storage VMotion, you can change the virtual disk format from thick to thin and reclaim this over-allocated but unused storage space (Figure 1.12).

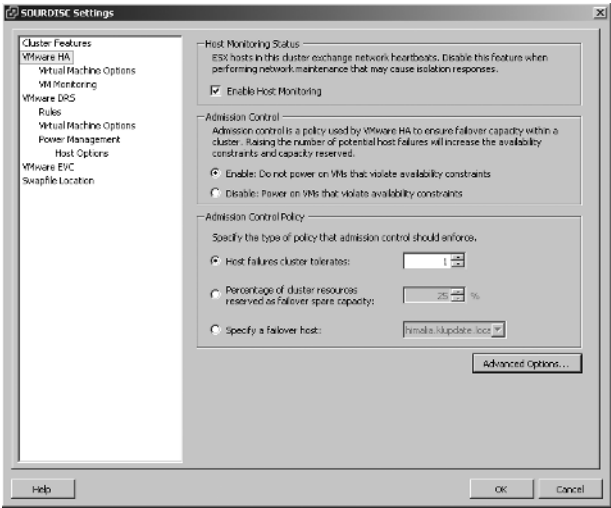
FIGURE 1.12
Storage VMotion,
virtual disk format
conversion



VMWARE HIGH AVAILABILITY

VMware High Availability (HA) extends high availability for any application running in a virtual machine, regardless of its operating system or underlying hardware configuration (Figure 1.13). VMware HA automatically detects physical server failures and operating system failures within virtual machines. If a failure is detected, VMware HA automatically restarts the affected virtual machine on a different ESX server in the resource pool. When restarting the affected virtual machines, you can also leverage DRS to take into account ESX host workload. The time it takes to reboot the virtual machine and restart the application is the only down-time that you will experience. VMware HA supports up to 32 ESX servers in a cluster. You can reserve a specific cluster capacity for failover or mark specific ESX servers as failover hosts. VMware HA can be a much simpler and cost-effective high-availability solution compared to traditional clustering solutions.


FIGURE 1.13
VMware HA
settings



VMWARE FAULT TOLERANCE

VMware Fault Tolerance (FT) provides zero downtime and zero data loss availability for all virtual machines during server hardware failures. When FT is enabled for a specific virtual machine, a secondary copy of that virtual machine runs in lockstep synchronization on another ESX host. This allows instantaneous, stateful failover between the two virtual machines and eliminates any disruption because of hardware failures on either host. FT does not support certain vSphere features and will display the warning shown in Figure 1.14 when you attempt to enable it.

FIGURE 1.14
VMware FT
warning



VMWARE DATA RECOVERY

VMware Data Recovery (VDR) provides a simple and easy-to-deploy disk-based backup and restore solution for protecting your virtual environment. VDR is deployed as a virtual appliance and does not require you to install any agents inside virtual machines. You back up and restore entire virtual machines and do not worry about the guest operating system or application running inside. It is integrated with VMware vCenter Server to provide a centralized user interface to manage your virtual machine backup and recovery jobs. VDR leverages virtual machine snapshots to make a precise copy in a short amount of time and also uses data *deduplication* technology to save on disk storage. VDR also supports Volume Shadow Copy Service (VSS) to enable consistent backups of virtual machines running Microsoft operating systems and applications. Depending on your business service-level agreements (SLAs) for Recovery Time Objective (RTO) and Recovery Point Objective (RPO), VDR can assist you in creating the perfect backup and recovery solution for your business organization. We discuss how you can use VDR in more detail in Chapter 2.

HIGH AVAILABILITY CLUSTERING WITH WINDOWS SERVER 2003 AND 2008

vSphere supports Microsoft Cluster Service (MSCS) with Windows 2003 and Failover Clustering for Windows 2008. Clustering is supported with both 32-bit and 64-bit guests. Booting from SAN virtual machines is supported as well. vSphere also supports Majority Node Set clusters with application-level replication, such as Exchange 2007 Cluster Continuous Replication (CCR). The serial attached SCSI (SAS) virtual device introduced as part of vSphere virtual hardware version 7 is required to support Failover Clustering configuration for Windows Server 2008 and R2 versions. Later chapters in this book will cover setting up Failover Clustering for Windows Server 2008 R2 and will demonstrate how to use this device.

TIP VMware provides a separate guide to detail the setup for Failover Clustering and Microsoft Cluster Service. The guide is available at www.vmware.com/pdf/vsphere4/r40_u1/vsp_40_u1_mscs.pdf.

Security

Security services help you secure your virtual infrastructure from vulnerabilities. They also enable applications to enforce an appropriate level of security policies in an operationally efficient way. VMware vSphere includes the following security services:

VMware vShield Zones VMware vShield Zones is an application-aware firewall that can be integrated with VMware vCenter Server to enforce corporate security policies and ensure regulatory compliance at the application level in vSphere environments. It continuously monitors and controls the network traffic flowing to and from virtual machines in its inventory, while still maintaining trust and network segmentation of users and sensitive data.

VMware VMsafe VMware VMsafe provides an application programming interface (API) that security vendors can leverage to develop VMware-aware security products. VMsafe enables partners to build virtual appliance-based security solutions that can monitor and protect virtual machine CPU state, memory pages, network traffic, and disk files and any processes executing inside them.

In the next sections, we will provide more details about each of these security services.

VMWARE vSHIELD ZONES

VMware vShield Zones helps you protect the privacy and confidentiality of virtual machines and their data. vShield Zones builds an inventory of the operating systems, applications, and open ports within your virtual data center. You can then use this information to monitor and enforce network access to sensitive areas of the virtual data center, including the DMZ, or to servers with sensitive data that is subject to regulations such as PCI, SEC 1742a, or SOX compliance. It also allows you to build logical trust or organizational boundaries within existing vCenter Server deployments, while still leveraging the flexibility and availability of shared resource pools. You can then define security policies to bridge, firewall, or isolate traffic across these boundaries.

The components that make up vShield Zones environment are as follows:

vShield This is a virtual appliance located on each vSphere host and is used to inspect traffic flow and provide firewall protection.

vShield Manager This manages all of the distributed vShield instances by providing monitoring, centralized configuration, and software updates.

Once deployed, vShield sits in between your protected virtual machines and the external network interfaces. This allows vShield to intercept each network packet and enforce the policies that have been created for the particular security zone.

VMWARE VMSAFE

VMware VMsafe enables an open security architecture with a set of APIs from VMware that gives security vendors the insight into the inherent properties of virtualization, similar to a hypervisor. Using this set of APIs, the security partners can develop virtual appliance-based security solutions that can monitor and protect virtual machine CPU state, memory pages, network traffic, and disk files and the processes executing inside them. Because these products will work in

conjunction with the virtualization layer, they are expected to provide higher levels of security to virtual machines compared to even physical servers. The following are examples of such VMsafe-based VMware-aware security products:

- ◆ A single antivirus virtual appliance that protects a vSphere host and all virtual machines running on it.
- ◆ A network security virtual appliance that protects network access for all virtual machines on a vSphere host.
- ◆ Security solutions that are aware of advanced VMware features such as DRS and vDS and that continue to protect virtual machines as they migrate using VMotion, Storage VMotion, or vDS network VMotion.

VMkernel PROTECTION

vSphere introduces several security mechanisms to assure the integrity of the VMkernel and loaded modules as they reside on disk and in memory. It leverages disk integrity mechanisms to protect against malware, which might attempt to overwrite or modify VMkernel as it persists on disk. vSphere makes use of Trusted Platform Module (TPM), a hardware device embedded in servers, to protect the boot-up of the hypervisor.

VMkernel modules are digitally signed and validated during load time to ensure the authenticity and integrity of dynamically loaded code. VMkernel also uses memory integrity techniques at load time coupled with microprocessor capabilities to protect itself from common buffer-overflow attacks that are used to exploit running code. All these techniques are part of ongoing efforts to protect the hypervisor from common attacks and exploits and create a stronger barrier of protection around the hypervisor.

Scalability

Scalability services allow the vertical and horizontal scaling of virtual infrastructure while ensuring that the right amount of resources are allocated to applications without any disruption.

VMware ESX and ESXi scalability vSphere continues to extend the maximum supported configurations for the underlying physical hardware used by ESX and ESXi as well as virtual machines running on them. vSphere also makes the switch to 64-bit versions of VMkernel and Service Console for better scalability.

VMware DRS VMware DRS improves scalability in vSphere deployments in two ways. First, it automatically and continuously balances the workload across ESX servers within a cluster, making sure that no single virtual machine is bottlenecked on resources. Second, it provides a proportional resource allocation mechanism using shares, reservations, and limits so that you can dynamically control the resources used by a virtual machine without the need for any reboot.

Virtual machine hot-add support Virtual hardware version 7 in vSphere 4 introduces hot-add support for various virtual devices. The ability to add more resources to a virtual machine without powering it off can help you improve virtual machine scalability as needed. This functionality is supported only if the underlying guest operating system supports it.

VMWARE ESX AND ESXi SCALABILITY

vSphere 4 extends the scalability of the virtualization platform in several aspects. vSphere 4 supports servers with up to 64 logical CPUs and 1 TB of RAM. Consequently, vSphere can support up to a total of 512 virtual CPUs per single host. By increasing the number of virtual machines that can run on single host, vSphere can achieve a higher consolidation ratio compared to earlier versions of ESX.

At an individual virtual machine level, vSphere 4 adds support for up to eight virtual CPUs and 255 GB of memory. With these higher per virtual machine scalability limits, you can now virtualize more enterprise applications in your data center than before.

vSphere has also switched to 64-bit versions of VMkernel and Service Console. The 64-bit versions offer scalability benefits over 32-bit versions similar to other software applications. These benefits include the ability to address a lot more memory and better support for the newer 64-bit hardware.

VMWARE DRS

We discussed VMware DRS in detail in the “Infrastructure Services” section. Here we’ll focus only on those DRS features that help improve scalability in vSphere deployments. DRS enables the ongoing dynamic load balancing of server resources within a cluster. This ensures that the applications are getting required resources all the time and no single virtual machine is bottlenecked on resources. The optimal virtual machine placement resulting from DRS load balancing can improve application scalability in a way that you cannot beat on a continuous basis.

DRS also allows you to allocate resources to virtual machines in proportion to the priorities established in the form of shares, reservations, and limits. For example, DRS will allocate proportionally higher shares of resources to virtual machines with higher shares or will guarantee fixed quantities of memory or CPU for virtual machines based on their reservations. You can dynamically shrink and grow virtual machine resource usage as needed without rebooting them.

Together, the dynamic load balancing and proportional resource allocation mechanisms can improve application scalability in vSphere deployments.

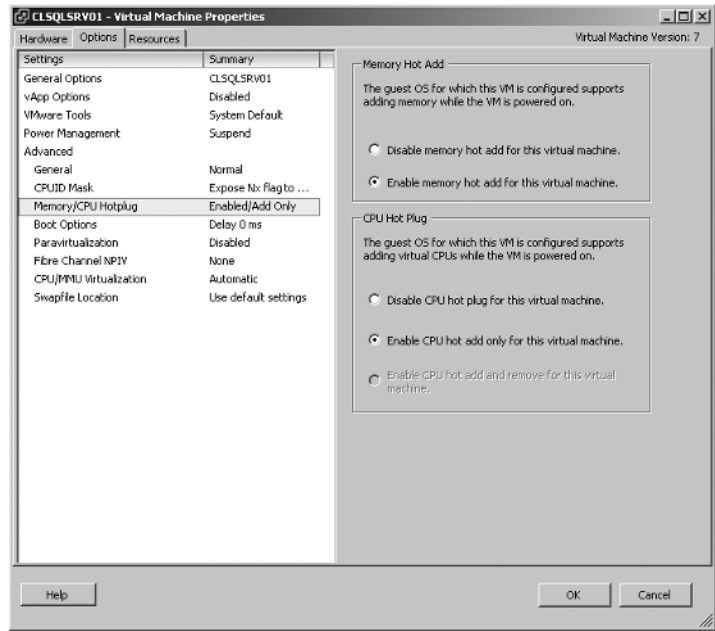
VIRTUAL MACHINE HOT-ADD SUPPORT

Virtual hardware version 7 introduced in vSphere 4 has hot-add support for various virtual devices:

- ◆ You can hot add CPU and memory to a virtual machine when needed without rebooting it (Figure 1.15).
- ◆ You can add or remove virtual storage and network devices from virtual machines without disruption.
- ◆ You can hot extend virtual disks for running virtual machines without any downtime.

This ability to dynamically add more resources to a virtual machine without powering it off can help you scale virtual machines as needed. However, this functionality needs to be supported by the guest operating system running inside the virtual machine. As a result, virtual machine hot-add support is currently supported for a limited number of operating systems only.

FIGURE 1.15
Virtual machine
hot-add support for
CPU and memory



Management Services

VMware vSphere management services include the tools and interfaces available to virtual infrastructure administrators to efficiently manage VMware vSphere environments. These management services include the following:

vCenter Server vCenter Server provides the central point of control for the administration and management of VMware vSphere environments. It acts as a single console to manage key IT processes for a virtualized data center such as monitoring, provisioning, patching, migration, capacity management, disaster recovery, and many other critical functions. VMware vCenter also integrates with existing industry-leading systems management tools from vendors such as BMC, CA, HP, and IBM/Tivoli for end-to-end physical and virtual management for a data center.

vCenter Orchestrator vCenter Orchestrator is a powerful orchestration engine integrated with vCenter. You can either use an out-of-the-box library of workflows or develop your own customized workflows to automate operational tasks.

vCenter Guided Consolidation vCenter Guided Consolidation provides a wizard-based process for first-time virtualization users to accelerate the server consolidation process. You can utilize this in smaller environments to discover physical servers, analyze their resource utilization, convert these physical servers to virtual machines, and place them on appropriate vSphere hosts.

vCenter Update Manager vCenter Update Manager automates the process of scanning and patching online VMware ESX hosts and selective Microsoft and Linux virtual machines to enforce compliance to patch standards. You can upgrade virtual machine hardware, VMware Tools, and virtual appliances as well as patch and update third-party software running on the virtual machines and virtual appliances.

vSphere Command-Line Interfaces vSphere provides two key command-line interfaces (CLIs) to automate common vSphere administration and management tasks: vSphere CLI and vSphere PowerCLI.

In the next sections, we will provide more details about each of these management services.

vCenter Server

VMware vCenter Server provides a single-pane-of-glass view for managing and administering the infrastructure and application services described in the previous sections. It provides visibility into every aspect of VMware vSphere environments and enables unified management of all the ESX/ESXi hosts and virtual machines in your data center from a single console along with aggregate performance monitoring of clusters, hosts, and virtual machines. VMware vCenter Server gives administrators insight into the status and configuration of clusters, hosts, virtual machines, storage, the guest OS, and other critical components of a virtual infrastructure—all from one place.

VMware vCenter Server lets you rapidly provision virtual machines and hosts using standardized templates. It creates a comprehensive map of the inventory and topology of your data center. It also controls access to virtual assets and functionality through fine-grained access controls, custom roles, permissions, and integration with existing Microsoft Active Directory authentication mechanisms. VMware vCenter Server also gives administrators control over key capabilities such as VMware VMotion, Distributed Resource Scheduler, High Availability, and Fault Tolerance.

In addition, VMware vCenter Server uses an open plug-in architecture to create a scalable and extensible platform. This allows VMware partners to integrate with vCenter Server to develop advanced management capabilities in areas such as capacity management, compliance management, business continuity, and storage monitoring. The vCenter Server APIs also allow customers to integrate physical and virtual management tools, using their choice of enterprise management tools to connect to vCenter Server.

vCenter Server 4 has been enhanced to include several new features such as the following:


- ◆ Simplified navigation and inventory search
- ◆ Host profiles
- ◆ Linked Mode
- ◆ Hardware monitoring with CIM
- ◆ Centralized licensing

We will provide more details about these features in the next sections.

SIMPLIFIED NAVIGATION AND INVENTORY SEARCH

VMware vCenter Server 4 has redesigned the landing page to provide users with an easy, one-click access to the vCenter inventory, vCenter solution plug-ins, and key administration and management tasks (Figure 1.16). In addition to the home page, vCenter Server also introduces a globally accessible navigation bar that makes it easy to switch between different parts of vSphere Client. The client application also remembers the view that was displayed when you log out of vSphere Client and will return you to that view when you next log in.

FIGURE 1.16
VMware vCenter
inventory search



vCenter Server 4 also includes a new inventory search field to simplify locating virtual machines, hosts, data stores, networks, and folders based on the criteria you specify. You can perform simple searches based on keywords entered in the search field at the top right of vSphere Client. You can also perform advanced searches by specifying multiple search criteria—for example, virtual machines that need an update to VM Tools or data stores that have less than 10 GB of free space remaining.

HOST PROFILES

vCenter Server 4 introduces host profiles that can simplify host configuration management through user-defined configuration policies (Figure 1.17). You can use profile policies to eliminate any manual, single-host, ad hoc ESX host configuration and efficiently maintain configuration consistency and correctness across the entire data center.

Using host profile policies, you capture the blueprint of a known, validated “golden” configuration and use this to quickly configure networking, storage settings, security settings, and so on, across a large population of hosts. For example, you can quickly update DNS and NTP settings for several ESX hosts or configure multiple ESX hosts to use a new vNetwork Distributed Switch. Note that ESX hosts need to be in maintenance mode to apply a host profile.

You can also use host profile policies to monitor and report on compliance to standard host configuration settings across the data center (Figure 1.18). The profile compliance information displayed on the Profile Compliance tab depends upon the object selected in the vCenter inventory panel. For example, if you select a cluster in the vCenter inventory panel, you will see profile compliance information for all hosts within the selected cluster. vCenter generates a base set

of host compliance checks based on host configuration settings, such as networking, DNS, and NTP settings. In addition, vCenter performs certain built-in cluster compliance checks (shown in Table 1.1) even if a host profile is not attached to a cluster.

FIGURE 1.17
Host profiles' NTP configuration

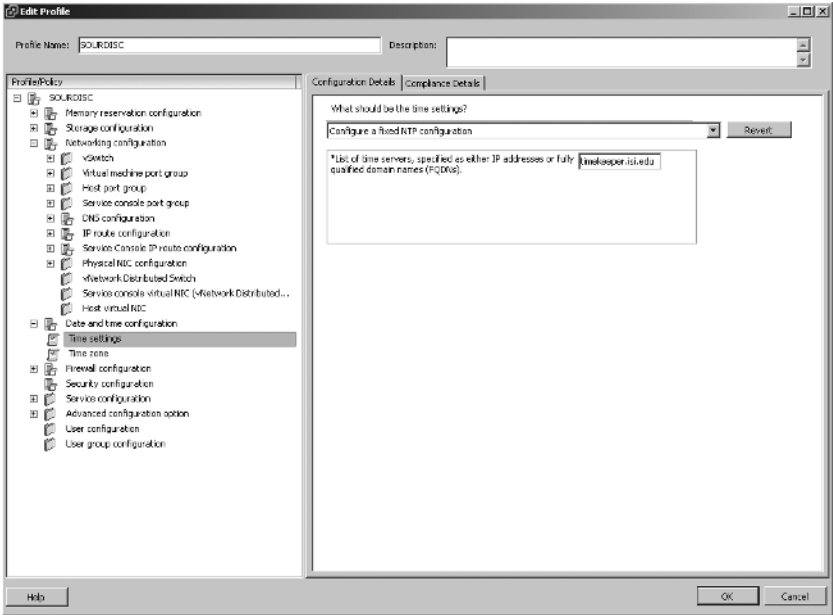


FIGURE 1.18
Profile compliance

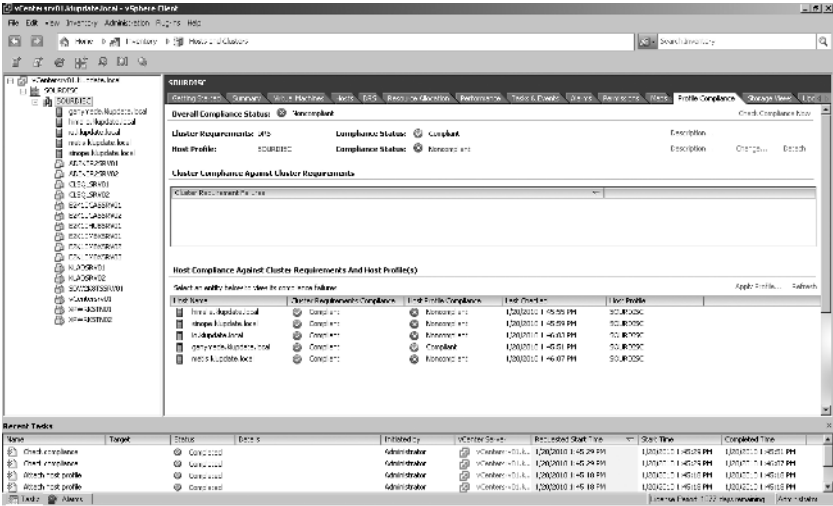


TABLE 1.1: Built-in Cluster Compliance Checks

CLUSTER REQUIREMENT	CLUSTER COMPLIANCE CHECK
VMware DRS	Validate that VMotion NIC speed is at least 1000 Mbps. Validate that VMotion is enabled. Validate that at least one shared data store exists.
VMware DPM	Validate that power management is supported on the host.
VMware HA/VMware FT	Validate that FT logging is enabled. Validate that FT logging NIC speed is at least 1000 Mbps. Validate that all the hosts in the cluster have the same build for FT. Validate that the host hardware supports FT.

vCENTER SERVER LINKED MODE

A single instance of vCenter Server 4 can manage up to 300 hosts and 3,000 virtual machines. This limit is not enough for organizations ramping up their virtualization projects to span an entire data center. To meet these higher scalability requirements, vCenter Server 4 introduces a feature called Linked Mode, where multiple vCenter Server systems can be linked together and monitored from a single vSphere Client session. Using the Linked Mode feature, you can manage up to 1,000 hosts and 10,000 virtual machines across 10 vCenter Server instances.

Linked Mode leverages Microsoft Active Directory Application Mode (ADAM), an implementation of Lightweight Directory Access Protocol (LDAP), to store and synchronize data across multiple vCenter Server instances. ADAM is installed automatically as part of the vCenter Server 4 installation. The ADAM instances in a group use peer-to-peer networking to replicate the following information for each vCenter instance to the LDAP directory:

- ◆ Connection information (IP and ports)
- ◆ Certificates and thumbprints
- ◆ Licensing information
- ◆ User roles and permissions

When vCenter Servers are connected in Linked Mode, you can do the following:

- ◆ Log in simultaneously to all vCenter Servers using a valid credential
- ◆ Search the inventories of all the vCenter Servers in the group
- ◆ View the inventories of all the vCenter Servers in the group in a single inventory view

When you log in vCenter using Linked Mode, you will see the inventory of all vCenter instances at once. The inventory tree on the left side will show each vCenter instance at the top level. You can then use the +/- indicator to expand or collapse the inventory tree to, respectively, show or hide lower-level objects such as the data stores, folders, clusters, hosts, and so on, for

any vCenter instance. Note that a user needs to have valid permissions to be able to see a vCenter Server instance.


The vCenter Linked Mode feature’s visibility across multiple vCenter instances applies only to view and search operations. Other operations are confined within a single vCenter inventory. For example, you cannot drag and drop a host between vCenter instances or a virtual machine between hosts on two different vCenter instances.

You can join a vCenter instance to a Linked Mode group at the time of installation or afterward by modifying an existing deployment. Both of these methods are described in the vSphere Installation Guide.

HARDWARE MONITORING WITH CIM

vCenter Sever 4 provides a Hardware Status plug-in to monitor and manage the health of your VMware ESX servers, including key components such as fans, system boards, and power supplies (Figure 1.19). The vCenter Hardware Status plug-in uses the industry-standard Common Information Model (CIM) interface to display this health information. It implements the System Management Architecture for Server Hardware (SMASH) profiles defined by the Distributed Management Task Force (DMTF).

FIGURE 1.19
vSphere hosts’
hardware status



The vCenter Hardware Status plug-in provides you with an integrated, centralized view of both the physical and virtual assets in your vSphere environment. The plug-in is also integrated into the vCenter alarm interface so you can be alerted when hardware failures occur. You can also trigger automated alarm workflows based on these vCenter alarms to preempt and remedy hardware problems. For example, if vCenter detects that the host temperature is getting too hot, it could trigger an alarm action that automatically puts the VMware ESX host into maintenance mode (migrating the VMs off the host using VMware DRS in the process) to allow the server to cool down.

CENTRALIZED LICENSING

In vSphere 4, VMware has redesigned the licensing facility. License reporting and management are centralized into vCenter Server 4. Unlike VMware Infrastructure 3, there is no separate license server that must be installed and monitored. Using the vCenter Server interface, you can centrally assign VMware vSphere licenses, report on license usage, and monitor for license state and compliance. If you upgrade all of your hosts, you no longer need a license server or host license files. All product and feature licenses use 25-character license keys that you can manage and monitor from vCenter Server 4.

When vCenter Server 4 assigns a license key to a vSphere host, the license key is copied to the host and saved in a persistent format. In the event that the host becomes disconnected from vCenter Server 4, the license key remains active on the host indefinitely, even after a host reboot. Only a deliberate licensing operation by the user can remove or replace a host license key. VMware recommends assigning all VMware vSphere licenses centrally from vCenter Server 4; however, you can also assign license keys directly to individual hosts.

For more information on VMware vSphere licensing, go to www.vmware.com/support/licensing.html.

vCenter Orchestrator

vCenter Server 4 includes vCenter Orchestrator, a workflow orchestration engine to help you automate management tasks for vSphere either by using out-of-the-box workflows or by assembling workflows using an easy drag-and-drop interface.

Workflows are reusable building blocks that combine actions, decisions, and results that, when performed in a particular order, complete a specific task or process in a virtual environment. You can leverage the out-of-the-box library of extensible workflows provided with vCenter Orchestrator to create and execute automated, configurable processes to manage your infrastructure. Or you can use vCenter Orchestrator to capture best practices within your data center for tasks, such as provisioning virtual machines, backing up, and performing regular maintenance, and turn them into workflows. vCenter Orchestrator also provides access to all operations in the vCenter Server API. You can integrate all these operations into your automated processes or workflows. This will help you achieve consistency, standardization, and overall compliance with existing IT policies in your virtual environment.

Orchestrator also allows you to integrate with other third-party management and administration solutions through its open plug-in architecture.

vCenter Guided Consolidation

vCenter Guided Consolidation provides a wizard-based process for first-time virtualization users to accelerate the server consolidation process. The Guided Consolidation service is provided as a modular plug-in to vCenter Server. To reduce the overhead on the vCenter Server, the Guided Consolidation service can also be installed on a different system than vCenter Server 4.

As the name suggests, Guided Consolidation guides you through all the steps of the consolidation process:

- ◆ Automatic discovery of physical servers
- ◆ Performance analysis to identify virtualization candidates

- ◆ Physical-to-virtual conversion
- ◆ Intelligent placement on the right vSphere host

You can find and analyze physical systems using Guided Consolidation in several ways:

- ◆ You can automatically scan active domains daily.
- ◆ You can manually add systems by entering a computer name, a singular or range of IP addresses, or a filename.

Guided Consolidation is recommended for smaller environments, up to 100 concurrent physical machines at any given time. For larger environments, you should use VMware Capacity Planner.

vCenter Update Manager

vCenter Update Manager automates the process of patch management in vSphere environments and helps you enforce compliance to patch standards. The vCenter Update Manager can help you do the following:

- ◆ Patch and upgrade VMware ESX/ESXi hosts
- ◆ Apply patches to Windows and certain versions of Linux guest operating systems in virtual machines
- ◆ Upgrade VMware Tools and virtual hardware for virtual machines
- ◆ Patch and upgrade virtual appliances

VMware Update Manager integrates with the Shavlik patch management technology to provide these patching capabilities. It also allows you to download patches from a remote server and stage them to a local server. Staging patches on local server can be very helpful when patching several hosts and virtual machines.

Update Manager also allows you to create update baselines and baseline groups that represent a set of updates. You can then review the Update Manager dashboard to compare hosts and virtual machines against these baselines. If the host or virtual machine is not compliant, then it can be easily remediated. You can use this to ensure consistency across ESX hosts and virtual machines in your vSphere environment.

In Chapter 2, we discuss in more detail how you can use Update Manager to patch Windows Server 2008 virtual machines and comply with your patch management policies.

vSphere Management Assistant

With vSphere, VMware also ships vSphere Management Assistant (vMA), a virtual machine that includes the vSphere CLI, vSphere SDK for Perl, and other prepackaged software. You can use this prebuilt virtual machine to run agents and scripts to manage ESX/ ESXi and vCenter Server systems. Some folks refer to vMA as the missing service console for ESXi. A key functionality of vMA includes noninteractive login, which allows you to use the vSphere CLI without having to explicitly authenticate each time. vMA can also collect ESX/ESXi and vCenter Server logs and store the information for analysis. vMA can also host third-party agents for added management functionality.

vSphere Host Update Utility

vSphere also includes the vSphere Host Update Utility to enable you to remotely upgrade ESX hosts (versions 3.0 and newer) to ESX 4.0. It upgrades the virtual machine kernel (VMkernel) and the Service Console, if present, but does not upgrade VMFS data stores or virtual machine guest operating systems. vSphere Host Update Utility comes as a stand-alone Microsoft Windows application that provides real-time status of a remote upgrade as well as allows you to specify custom post-upgrade scripts. The utility performs an automated host compatibility check as a pre-upgrade step to verify that each host is compatible with ESX 4.0/ESXi 4.0. The Host Update utility also includes a new rollback feature to back out failed upgrades. It is recommended for smaller deployments with fewer than 10 ESX/ESXi hosts, without vCenter Server or Update Manager.

vSphere Command-Line Interfaces: vCLI and PowerCLI

vSphere provides two key command-line interfaces to automate vSphere administration and management tasks:

vSphere CLI The VMware vSphere CLI (vCLI) provides command-line interface tools for managing vSphere ESX and ESXi servers. You can use this with ESX/ESXi and vCenter Server 4 systems as well as with ESX/ESXi versions 3.5 Update 2 and newer. vCLI includes more than 30 command-line utilities to help you provision, configure, and maintain your ESX and ESXi hosts.

vSphere PowerCLI Based on Microsoft PowerShell technology, VMware vSphere PowerCLI is a powerful command-line tool that lets you automate all aspects of vSphere management, including network, storage, ESX, VM, guest OS, and more. PowerCLI is distributed as a Windows PowerShell snap-in, and it includes more than 150 PowerShell cmdlets, along with documentation and samples.

VMware vApp

VMware vSphere includes support for vApp, a single logical entity that encapsulates an *n*-tier application comprising one or more virtual machines. vApp uses the industry-standard Open Virtualization Format (OVF) to specify and encapsulate all components of a multitier application as well as the operational policies and service levels associated with it. Along with virtual machines, vApps also captures the dependencies between these virtual machines as well as the resource allocations associated with these virtual machines. Defining all this information in one logical entity makes it very convenient, often in a single step, to power off/on, clone, deploy, and monitor an entire application. vCenter Server can create and run vApps, as well as import and export them in compliance with the OVF 1.0 standard.

VMware Management Solutions

Complementing vCenter Server and its management services, VMware also provides several virtualization management solutions. Here are a few of these products:

- ◆ VMware vCenter AppSpeed to monitor and troubleshoot the performance of applications virtualized in a VMware vSphere environment
- ◆ VMware vCenter CapacityIQ to perform capacity sizing and management of virtual machines, clusters, and data centers

- ◆ VMware vCenter Chargeback to automate tracking and chargeback for the cost of IT services
- ◆ VMware vCenter Lab Manager to create an on-demand lab infrastructure
- ◆ VMware vCenter Lifecycle Manager to automate and manage virtual machine provisioning in your virtual infrastructure
- ◆ VMware vCenter Site Recovery Manager to automate the setup, testing, and failover of disaster recovery processes

In this section, we provide a brief overview for each of these solutions.

VMware vCenter AppSpeed

VMware vCenter AppSpeed enables you to monitor application-level latency, throughput, and transaction rate and correlate the data with performance indicators at the infrastructure level. It can monitor most multitier applications, including database traffic (Microsoft SQL, MySQL, Oracle), any HTTP or HTTPS traffic (most application servers and web servers), and Microsoft Exchange Server traffic. It does this without installing an agent or requiring login credentials to the applications. AppSpeed is shipped as a virtual appliance, so it can be configured and deployed with minimal effort. When deployed, AppSpeed monitors application traffic through virtual switches and automatically builds application topology and maps interdependencies between application components. AppSpeed can detect the slowest or most used application transactions. It also allows you to drill down into individual transactions, as well as point out transaction dependencies across components in a multitier application stack. This can help you quickly isolate and remediate application performance problems in vSphere deployments. With AppSpeed, you can also set up ongoing performance service-level monitoring based on thresholds for transaction latency. You can either set this threshold manually or let AppSpeed automatically select the values based on the historical information it has.

We discuss how you can deploy and use AppSpeed in Chapter 5. For more details, you can also check the vCenter AppSpeed product page at www.vmware.com/products/vcenter-appspeed/.

VMware vCenter CapacityIQ

VMware vCenter CapacityIQ adds capacity management capabilities to your vSphere deployment. It enables ongoing, real-time tracking of capacity utilization in your environment and analyzes this information to perform capacity trending and forecasting. With CapacityIQ, you can tell how long your capacity will last and when it's time to add more capacity. CapacityIQ can also help you optimize current capacity usage and size your infrastructure correctly by identifying idle, powered-off, over-provisioned, and under-provisioned virtual machines. CapacityIQ also provides "what-if" modeling to help you understand the impact of adding or removing virtual machines or vSphere hosts to your deployment. You can model the virtual machines or vSphere hosts based on expected usage or even base them on existing VMs. CapacityIQ is shipped as a virtual appliance, so it can be configured and deployed with minimal effort.

For more details, check the vCenter CapacityIQ product page at www.vmware.com/products/vcenter-capacityiq/.

VMware vCenter Chargeback

VMware vCenter Chargeback enables you to track how virtual infrastructure is being used by different departments in your organization, correlate it to their cost centers, and finally charge them for the consumption. Chargeback lets you create variety of cost models—fixed, reservation, usage, or a hybrid cost model—that measure cost on a per-VM basis. These cost models can be applied at the business level in a chargeback hierarchy, so there is no need to create cost models for each VM. It supports allocation-based or utilization-based chargeback, or a hybrid of both. You can leverage this to set up tiered cost models that vary across cost centers, or types of use requirements. For example, a production environment may cost more than a development environment. Similarly, you will need to pay a premium for a highly available virtual machine with HA or FT enabled. Using Chargeback's detailed reports, you can either bill or simply "show" to the business groups the IT cost of their virtual infrastructure usage.

For more details, check the vCenter Chargeback product page at www.vmware.com/products/vcenter-chargeback/.

VMware vCenter Lab Manager

VMware vCenter Lab Manager allows you to deploy multiple virtual machines for a development or test environment from a library of existing configurations using a self-service portal. It has an intuitive user interface that is simple enough for non-IT users but also includes advanced features such as system quotas and access control to manage resource allocation. You can control multiple machines as a single atomic unit and deploy them to a development or test or staging environment. It is also widely used in training and education to provision lab environments, in support and help-desk groups for troubleshooting, and in sales and marketing for live product demos and for product evaluations.

For more details, check the vCenter Lab Manager product page at www.vmware.com/products/labmanager/.

VMware vCenter Lifecycle Manager

VMware vCenter Lifecycle Manager provides you with better control over virtual machine provisioning and deployment in your vSphere deployment. It provides workflows built on top of vCenter Orchestrator to automate the entire virtual machine life-cycle management from provisioning to decommissioning. Lifecycle Manager provides a centralized web portal to enable self-service provisioning of virtual machines. The end users choose from a catalog of virtual machine configurations and are allowed to request or create virtual machines based on user policies and access control rights. Lifecycle Manager then automatically tracks the virtual machine throughout its life.

For more details, check the vCenter Lifecycle Manager product page at www.vmware.com/products/lifecycle-manager/.

VMware vCenter Site Recovery Manager

VMware vCenter Site Recovery Manager addresses the disaster recovery challenges for the entire data center. By leveraging virtualization and replication capabilities of storage, it enables more affordable disaster recovery protection than previously possible. With vCenter Site Recovery Manager, you can build multiple disaster recovery workflows within your vCenter

implementation to automate your failover plans. You can also test these workflows completely so that you know the recovery will work if and when you need it. It utilizes storage replication to ensure that data is successfully and accurately transferred to the failover site. Because of virtualization, you do not have to maintain strict hardware compatibility or map servers one to one with your primary site environment.

For more details, check the vCenter Site Recovery Manager product page at www.vmware.com/products/site-recovery-manager/.

VMware vSphere Editions

VMware provides vSphere in two separate offerings: one targeted for small businesses and the other for midsize and enterprise businesses. Each of these offerings includes different editions of vSphere that are tiered, based on the features included. In this section, we will provide a quick overview of these offerings.

vSphere for Small Businesses

For the small businesses, these are the vSphere editions available:

- ◆ VMware vSphere Essentials
- ◆ VMware vSphere Essentials Plus

Both these editions are limited to three physical servers and allow you to create virtual machines with up to four virtual CPUs. Unlike the vSphere editions for midsize and enterprise businesses, these editions also include the license for VMware vCenter Server for Essentials to help you manage the virtual environment. The Essentials Plus edition adds VMware HA and VMware Data Recovery to the Essential edition to provide the high-availability and data protection features. Note that these editions for small businesses are self-contained packages that cannot be combined with other VMware vSphere editions.

If you want to virtualize only one physical server, you can use VMware ESXi, which is the free version of the VMware ESX hypervisor with a smaller footprint. You will need to upgrade to other vSphere editions to leverage any additional vSphere features or use vCenter Server for management.

To address the needs of smaller deployments at retail and branch offices, VMware provides special editions of the vSphere Essentials packages:

- ◆ VMware vSphere Essentials for Retail and Branch Offices
- ◆ VMware vSphere Essentials Plus for Retail and Branch Offices

These retail and branch office editions allow you to upgrade the vCenter Server for Essentials edition to vCenter Server Standard edition. This Standard edition adds support for vCenter Orchestrator and vCenter Linked Mode features. The Essentials Plus edition adds VMware HA and VMware Data Recovery to the Essential edition to provide the high-availability and data protection features.

You can find the detailed comparison between these small business editions at www.vmware.com/products/vsphere/buy/small_business_editions_comparison.html.

vSphere for Midsize and Enterprise Businesses

VMware provides the following vSphere editions for mid-size and enterprise businesses:

- ◆ VMware vSphere Standard
- ◆ VMware vSphere Advanced
- ◆ VMware vSphere Enterprise
- ◆ VMware vSphere Enterprise Plus

These editions are tiered based on the vSphere features included. You can find the detailed comparison between these midsize and enterprise business editions at www.vmware.com/products/vsphere/buy/editions_comparison.html.

VMware vCenter Server is a required component for vSphere deployment and has to be purchased separately for these midsize and enterprise editions of vSphere. vCenter Server is also available in two editions:

- ◆ VMware vCenter Server Foundation
- ◆ VMware vCenter Server Standard

The Foundation edition is limited to managing three vSphere hosts. The Standard edition removes that limit and also adds support for vCenter Orchestrator and vCenter Linked Mode features.

TIP VMware provides a vSphere Purchase Advisor that can help you figure out the right edition of vSphere for your environment. The tool is available at www.vmware.com/products/vsphere/purchase-advisor/.

vSphere Compatibility Guides

VMware maintains strict compatibility lists for its virtualization platform. This is one of the key reasons why vSphere is rated as a mainframe-like, highly reliable, and resilient virtualization platform. vSphere is certified across the complete stack of servers, storage, operating systems, and software applications. VMware works closely with the various hardware providers to certify new hardware with vSphere. You can check the compatibility guides at www.vmware.com/resources/compatibility. Because of ongoing certifications, these compatibility guides are updated from time to time. You can leverage this information to plan and deploy hardware and operating systems across the entire data center. The online compatibility guides allow you to search for vSphere compatibility in the following categories:

- ◆ Server systems
- ◆ Storage/SAN
- ◆ I/O devices
- ◆ Guest/host operating systems
- ◆ VMware View-compatible devices

The systems, storage, and I/O devices compatibility guides list the hardware that's compatible with various versions of VMware ESX and ESXi. The guest and host operating system compatibility guides list the operating systems supported by various VMware products such as ESX/ESXi, Workstation, Fusion, and Server.

Summary

With the release of vSphere, VMware continues to expand the capabilities of its successful VMware Infrastructure virtualization platform. VMware classifies the vSphere platform functionality into three groups of services:

- ◆ Infrastructure services
- ◆ Application services
- ◆ Management services

Each of these services further comprises several components that together make vSphere a reliable and scalable virtualization platform. Although infrastructure and application services are part of vSphere, the management services are provided by VMware vCenter Server. In this chapter, we provided a high-level overview of these services and details on the different components that make these services. Understanding the capabilities of the underlying virtualization platform is the key to subsequent successful application deployment. You can leverage this information when we explain how to successfully virtualize the tier-1 Microsoft server applications in the next chapters.